

ProLiant BL p-Class GbE Interconnect Switch Overview

white paper



Abstract.....	3
Introduction.....	3
GbE Interconnect Kit	4
Modular hot-swappable design.....	4
Support for copper and fiber-based networks	4
Interconnect switch architecture	4
Downlink ports.....	5
LAN interconnect module uplink ports.....	5
Front panel ports.....	6
Crosslink ports	6
Maximizing network cable reduction	7
Network load balancing.....	8
Investment protection.....	8
Switch management.....	8
Rapid deployment, back-up, and restore	9
IP addressing.....	10
Web browser-based interface.....	10
Menu-driven console and scriptable command line interface	11
Simple network management protocol and remote monitoring.....	11
Switch security.....	12
Serviceability and diagnostics	12
Virtual LAN	13
Spanning tree.....	14
EtherChannel compatible link aggregation.....	14
Internet Group Management Protocol	14
Data Storm Prevention	14
Quality of Service	14

Enterprise-class performance	14
Network adapter teaming.....	15
Conclusion.....	15
Appendix: Summary of features of the ProLiant BL p-Class GbE Interconnect Kit.....	16
Switch performance.....	16
Switch network features	16
Switch deployment and configuration	16
Switch diagnostics and monitoring	17
Switch security.....	18
Switch availability	18
Switch ports per server blade enclosure	18
Switch physical and environmental.....	18
For more information.....	19

Abstract

This white paper provides an overview of the first generation Gigabit Ethernet (GbE) Interconnect Kit options for the ProLiant BL p-Class system. Each GbE Interconnect Kit includes a pair of industry-standard Ethernet switches that dramatically reduce the number of Ethernet network cables attached to the rear of the ProLiant BL p-Class server blade enclosure.

The intended audience for this paper includes engineers and system administrators familiar with the overall network design of the ProLiant BL p-Class system. For readers not familiar with this overall network design, please see the [ProLiant BL p-Class Networking Overview](#)¹ white paper.

Introduction

The HP ProLiant BL p-Class system includes a portfolio of ProLiant BL p-Class server blades, the p-Class server blade enclosure, and four interconnect options.

The ProLiant BL p-Class server blade family includes the dual-processor ProLiant BL20p and BL20p G2 server blades for mid-tier applications and the quad-processor ProLiant BL40p server blade for back-end applications.

The 6U (10.5 inch) p-Class server blade enclosure has eight server bays, and two outermost interconnect bays for collection of the network and storage signals. The signals are routed from the server blades, across the server blade enclosure backplane, and to the interconnect blades.

HP offers a family of interconnect options so that customers can choose how the Ethernet network signals exit the server blade enclosure. Interconnect options include two patch panel pass-through kits and two integrated Ethernet switch kits (the GbE Interconnect Kit and the GbE2 Interconnect Kit). The two patch panel options allow all Ethernet network signals to pass through to third-party LAN devices, thus giving customers flexibility in choosing their own switches. The GbE and GbE2 Interconnect Kit options include managed Ethernet switching technology that provides up to a 32-to-1 reduction in network cables at the back of the BL p-Class enclosure. This cable reduction significantly reduces the time required to deploy, manage, and service ProLiant BL p-Class systems.

Please see the *ProLiant BL p-Class Networking Overview* white paper for a summary of all four interconnect options along with an interconnect decision chart to assist in choosing the appropriate type of interconnect switch to use.


This white paper describes the GbE Interconnect Switch. The GbE Interconnect Switch is intended for applications that require 100 megabits per second (Mb/s) Fast Ethernet network adapter (NIC) consolidation to Gigabit Ethernet uplinks.

¹ Available at <http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html>.

GbE Interconnect Kit

The GbE Interconnect Kit contains two redundant, hot-swappable GbE Interconnect Switches and two 4-port LAN interconnect modules. Users first insert the LAN interconnect modules into the top-left and top-right interconnect bays at the rear of the server blade enclosure and then install the interconnect switches into the interconnect bays from the front of the enclosure (Figure 1). The GbE Interconnect Switches consolidate the 32 Ethernet ports from the server blades into one to twelve external ports: four Gigabit Ethernet uplink ports (uplinks) and four 10/100T uplinks on the two LAN interconnect modules and four 10/100T ports on the front of the switches. If desired, users can configure the GbE Interconnect Kit to concentrate the 32 Ethernet signals into a single external Ethernet port (see the section titled "Maximizing network cable reduction").

Figure 1. Installation of LAN interconnect modules and GbE Interconnect Switches



Modular hot-swappable design

To enhance the system serviceability and simplify upgrades, both the GbE Interconnect Switch and the LAN interconnect modules are hot swappable. This modular design provides two key benefits:

- The GbE Interconnect Switches can be quickly and easily removed and replaced from the front of the rack without the need for re-cabling.
- Each GbE Interconnect Switch can support copper or fiber Ethernet network uplinks by using different LAN interconnect modules.

Support for copper and fiber-based networks

Two interconnect kit options are available: the C-GbE Interconnect Kit for copper-based networks and the F-GbE Interconnect Kit for fiber-based networks. These kits are identical with exception of the interconnect modules, which are described in the "Interconnect switch architecture" section.

Interconnect switch architecture


Each GbE Interconnect Kit provides an end-to-end, fully redundant architecture that maximizes network availability (Figure 2). Redundant network adapters are routed from each server blade bay to each hot-swappable interconnect switch (four NICs total per server bay) creating a fully meshed topology to the external Ethernet network. The interconnect switches are cooled by on-board redundant fans and they are powered by two independent feeds via the redundant, N+N hot-plug ProLiant BL p-Class power enclosure.

Users can independently disable or enable the switch ports as desired. Auto-MDI/MDIX with auto-negotiation of speed and duplex mode is supported. Each GbE Interconnect Switch includes the following Ethernet ports:

- Sixteen dedicated internal 10/100 Mb/s Ethernet downlink ports connecting the server NICs to the switch.
- Six external Ethernet ports for data, interconnect switch A and B management, Integrated Lights-out (iLO) management, and/or pre-boot execution environment (PXE) remote configuration.
 - Two 10/100/1000T and two 10/100T uplink ports (C-GbE) or two 1000SX and two 10/100T uplink ports (F-GbE) on the rear-mounted LAN interconnect module.
 - Two 10/100T ports on the switch front panel.
- Two dedicated internal 10/100 Mb/s Ethernet crosslink ports bundled as a multi-port trunk for switch-to-switch communication and failover, if desired.

These Ethernet ports are described in the following subsections.

Figure 2. ProLiant BL p-Class GbE Interconnect Kit architecture



Downlink ports



Each GbE Interconnect Switch includes sixteen pre-assigned, internal 10/100 Mb/s Ethernet downlink ports connecting the server blade network adapters signals to the switch. The signals are routed as Ethernet from the server blades, across individual category 5e (CAT5e) specified signal traces on the passive backplane assembly of the server blade enclosure, and then to the switches.

LAN interconnect module uplink ports

The LAN interconnect module uplink ports are typically used to connect the interconnect switches to the network infrastructure; however, they are standard Ethernet switch ports and may be used as desired. The GbE interconnect kits are identical with exception of the LAN interconnect modules. The C-GbE Interconnect Kit includes two QuadT interconnect modules, each with two 10/100/1000T

and two 10/100T Ethernet ports all with RJ-45 connectors. The F-GbE Interconnect Kit includes two DualTSX interconnect modules, each with two 1000SX Ethernet ports with LC connectors and two 10/100T Ethernet ports with RJ-45 connectors. Each interconnect module has an independent guide pin to ensure insertion alignment and a robust connector assembly to prevent it from being used with incompatible devices.

Table 1. Interconnect modules used with the C-GbE and F-GbE Interconnect Kits

Components	C-GbE Interconnect Kit	F-GbE Interconnect Kit
Interconnect switches	Two GbE Interconnect Switches	Two GbE Interconnect Switches
Interconnect modules	<p>Two QuadT interconnect modules with:</p> <ul style="list-style-type: none"> Two 10/100/1000 BASE-T/TX/T ports Two 10/100 BASE-T/TX ports Four RJ-45 connectors with link speed and activity LEDs. 	<p>Two DualTSX interconnect modules with:</p> <ul style="list-style-type: none"> Two 1000 BASE-SX ports Two 10/100 BASE-T/TX ports Two LC connectors and two RJ-45 connectors with link speed and activity LEDs 

Front panel ports

Two 10/100T external ports are conveniently located on the front panel. These front panel ports allow users to perform local management and diagnostic tasks without unplugging a dedicated uplink. This prevents disruption the production network to perform these tasks. See the “Serviceability and diagnostics” section for more details on the switch front panel.

When the front panel ports are not used for local management, one or both ports may be used as additional uplinks. A channel is provided at bottom of each interconnect switch to allow network cables to be routed from the front to the rear of the rack, or vice versa.

Crosslink ports

The two GbE Interconnect Switches are connected through a pair of redundant, 10/100 Mb/s Fast Ethernet crosslink connections that are bundled in a 200 Mb/s Cisco EtherChannel compatible multi-port trunk. The signals are routed as Ethernet from switch to switch via individual CAT5e specified signal traces on the passive backplane assembly of the server blade enclosure. The crosslinks permit communication between the switches, if desired. They also provide the ability to:

- Manage both switches, perform PXE server boots, and access all iLO interfaces using any number of external switch ports. Therefore, a single switch uplink can be used to perform all Ethernet management tasks.
- Configure the p-Class system for advanced ProLiant Network Adapter Teaming including switch fault tolerance. See the section “ProLiant Network Adapter Teaming” for additional information.
- Communicate with any server network adapter from any switch uplink port. As a result, any single uplink port may be used to communicate to all 32 NICs. This permits the 32-to-1 cable reduction and provides additional system redundancy. If uplinks on one switch or the connection to a switch were to fail, all NICs still can be accessed via the other switch.

Maximizing network cable reduction


For maximum (97 percent) cable reduction, the 32 Ethernet signals within the server blade enclosure can be concentrated into any one single external Ethernet port. This results in a total of seven Ethernet connections for a fully configured 42U rack of six server blade enclosures containing 192 network adapters.

Applications that utilize a single uplink port include testing and evaluation systems, server blade enclosures with a few installed servers, and applications that require minimal bandwidth. On a heavily utilized system, using a single uplink port for all 32 network adapters can cause a traffic bottleneck. For example, using one uplink from interconnect switch A requires the traffic from all the network adapters routed to switch B to travel over the two crosslinks (a 200 Mb/s path), previously shown in Figure 2. The crosslinks are intended primarily as a failover route and generally are not used as a primary path. For more optimal performance, at least one uplink port on each interconnect switch would be used. However, system administrators may use any combination from one to all twelve external Ethernet ports to increase bandwidth, to separate network and management data onto physically isolated ports, or to add redundant connections to the Ethernet network backbone.

Another means to achieve network cable reduction is to link the GbE Interconnect Switches with other ProLiant BL e-Class and p-Class interconnect switches from different server blade enclosures. This is ideal for customers with multiple blade enclosures within a rack. It also allows the network administrator to define the desired level of network blocking or oversubscription.

For example, Figure 3 shows a configuration with three fully populated BL p-Class server blade enclosures, each with the C-GbE Interconnect Kit installed. The interconnect switches are linked, or daisy chained, together in redundant configuration using the four gigabit uplinks that, in turn, are connected to the Ethernet network backbone. Each enclosure contains thirty-two 10/100 Ethernet network adapters with an aggregate bandwidth of 3.2 gigabits per second (Gb/s), for a total bandwidth of 9.6 Gb/s for entire system of three enclosures. However, since uplinks are daisy changed together, the maximum system throughput for this configuration is 4.0 Gb/s (the combined throughput of the four daisy chained gigabit ports). This configuration creates a 2.4x blocking ratio (9.6 Gb/s versus 4.0 Gb/s); nevertheless, it reduces the total Ethernet network cables at the rack level from 96 to 4, a 96 percent cable reduction, while maintaining redundant connections to the Ethernet network backbone.

Figure 3. ProLiant BL p-Class GbE Interconnect Switch linking



In this example, four gigabit uplinks per server blade enclosure were used. Because the two GbE Interconnect Switches provide a total of twelve external Ethernet ports per enclosure, a wide variety of linking configurations are possible. Additionally, the two interconnect switches within each enclosure are connected by redundant crosslinks, which provide additional configuration possibilities.

Network load balancing

ProLiant BL systems configured with the interconnect switch support three network load balancing solutions. Options exist for providing this functionality integrated within or exterior to the server blade enclosure.

For load balancing within the server blade enclosure, the F5 Network BIG-IP Blade Controller (BIG-IP) may be used. BIG-IP is a software option for ProLiant BL e-Class and p-Class systems that provides a very economical solution to load balancing and traffic management between multiple server blades that reside in a single to multiple server blade enclosures.

BIG-IP is available from F5 Networks as a software option installed on ProLiant server blades. One license is installed on a server making the server blade into a “dedicated” load balancer; no additional software can be installed on the server blade. For a redundant solution, two copies of BIG-IP are installed on two server blades. The server blade(s) with BIG-IP installed may perform load balancing on blade servers that reside in the same or different server blade enclosures, both e-Class and p-Class. The blade servers may be located anywhere on the network as long as load balanced traffic to and from the servers pass through the BIG-IP Blade Controller. For more information on BIG-IP Blade Controller for ProLiant BL systems, see <http://h71028.www7.hp.com/enterprise/html/4557-0-0-0-121.html>.

For load balancing exterior to the server blade enclosure, a third party layer 7 Ethernet switch or network load balancing appliance may be used. This traditional approach uses a multi-tiered architecture where the interconnect switches are connected to one or more layer 7 switches or network load balancing appliances. Layer 7 switches and network load balancing appliances are available from several network vendors including Cisco, F5, Nortel, and others. This solution is supported on both e-Class and p-Class configured with any interconnect kit.

Investment protection

The GbE Interconnect Kits are fully supported in the p-Class sever blade enclosure with any combination of ProLiant BL20p, BL20p G2, and BL40p server blades. Users can upgrade to the GbE Interconnect Kit from existing RJ-45 Patch Panel and RJ-45 Patch Panel 2 Kits without powering down the server blade enclosure or server blades.


Switch management

The GbE Interconnect Switch is an industry-standard managed Ethernet switch, which means that users configure and manage the switch like other industry-standard Ethernet switches. To aid users during initial deployment, the GbE Interconnect Switch includes a default configuration that is fully operational at initial boot.

A web browser-based interface (BBI), menu driven console interface (MDI), and a command line interface (CLI) with scripting capability are pre-installed in the switch firmware to configure, manage, and monitor the interconnect switch. The interconnect switch also supports Telnet access, simple network management protocol (SNMP), and remote monitoring (RMON). Any combination of the downlinks, crosslinks, and external ports can be disabled, enabled, configured, and monitored on a per port basis. Out-of-band and in-band access to the switch management interfaces are supported

locally and remotely from anywhere on the network. Administration of both GbE Interconnect Switches is possible through any uplink port, the serial port, or the two 10/100T ports conveniently located on the front panel of each switch (Figure 4).

Figure 4. ProLiant BL p-Class GbE Interconnect Switch front panel



Rapid deployment, back-up, and restore

HP Rapid Deployment Pack for Windows version 1.40 or greater (RDP for Windows) introduces a new feature called server-side scripting. With server-side scripting, interconnect switch scripts can be integrated in an RDP for Windows job for deployment of both blade servers and switches. This is ideal for using RDP for Windows to deploy a blade server and then configure associated switch VLANs; however, any scriptable interconnect switch parameter can be integrated.²

HP Rapid Deployment Pack for Windows version 1.40 includes Altiris Deployment Solution version 5.6 (Altiris v5.6). Altiris v5.6 provides a "Switch Add-on" utility that identifies and manages VLAN settings on the GbE Interconnect Switches. A utility's Windows 32-bit graphical interface displays a tabular view of the all switches that have been identified, along with their respective switch ports, switch description, and VLAN assignments.

Using either the supported Windows graphical interface or the command line interface, an administrator can:

- View the GbE Interconnect Switch MIB (Management Information Base) system information via SNMP version 1.
- View the GbE Interconnect Switch ports and the VLAN setting for each port.
- Add/delete a GbE Interconnect Switch device to/from the database.
- Modify the VLAN setting of a GbE Interconnect Switch port via script configuration changes using the CLI or manually using the Windows graphical interface.
- Assign physical connectivity of a discovered device to a switch port (Windows interface only).

² For more information, see the white paper *Using ProLiant Essentials Rapid Deployment Pack for scripted blade-based switch configuration* at <http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html>.

The GbE Interconnect Switch supports trivial file transfer protocol (TFTP) allowing a copy of the interconnect switch configuration file to be saved and downloaded either to the original switch or to a different interconnect switch. This provides a method to rapidly deploy multiple systems with similar configurations and to provide backup and restore capabilities. Configuration settings can be modified through the user interfaces or directly within the configuration file. The configuration file has an extensible markup language (XML) format, which allows it to be directly viewed, printed, and edited. The configuration file may even be reset to the factory default settings at any time.

Users can perform firmware operating system upgrades by using TFTP through any external Ethernet port after boot-up, and by using ZModem through the serial interface during boot-up. The interconnect switch simplifies system upgrades by retaining its configuration after a firmware upgrade and by supporting the HP Support Paq automated firmware upgrade process for Windows deployment stations.

IP addressing

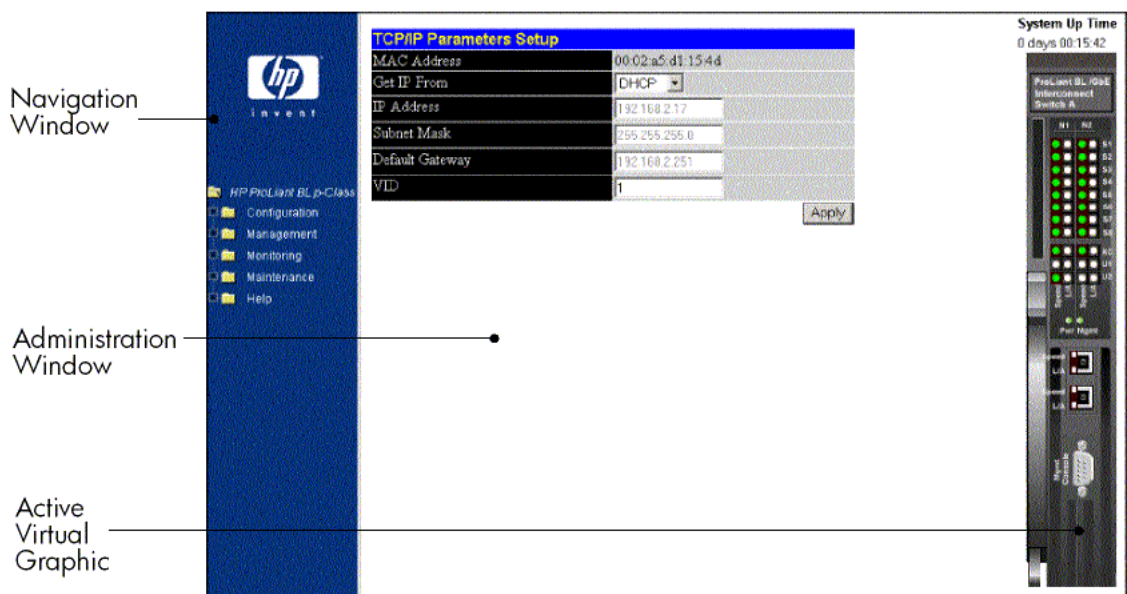
Users can configure the GbE Interconnect Switch to automatically obtain an IP address from a dynamic host configuration protocol (DHCP) or bootstrap protocol (BOOTP) server or they can manually assign an IP address through the command line and menu-driven console interfaces. Users can also assign an IP address from the BBI; however, they would have to reconnect with the newly assigned IP address. For increased security, an administrator can specify the IP-based management stations that are allowed to access the interconnect switch.

Web browser-based interface

Users can access the BBI by using Internet Explorer or Netscape Navigator over a TCP/IP network (Figure 5). The BBI interface consists of three main sections:

- The Active Virtual Graphic provides real time status of the switch front panel and provides the means to quickly view statistics of individual ports.
- The Navigation Window contains particular items or features to select.
- The Administration Window contains options for viewing or altering GbE Interconnect Switch information.

Figure 5. Browser-based interface for ProLiant BL p-Class GbE Interconnect Switch



Menu-driven console and scriptable command line interface

The menu-driven console and the command line interfaces provide the same configuration options as the BBI, plus the added benefits of an out-of-band connection and CLI scripting (Figure 6).

There are three methods of accessing these interfaces:

- Locally via the front panel RS-232 console port on each interconnect switch.
- Remotely using serial line IP (SLIP) via the front panel RS-232 console port on each interconnect switch.
- Remotely using a console Telnet session via any interconnect switch external Ethernet port.

Figure 6. Command line interface for ProLiant BL p-Class GbE Interconnect Switch



Simple network management protocol and remote monitoring

The GbE Interconnect Switch supports industry-standard SNMP management information bases (MIBs), HP enterprise switch MIBs, environmental traps, and RMON groups 1 (statistics), 2 (history), 3 (alarms), and 9 (events). The SNMP agents are pre-installed in the interconnect switch firmware. Four community strings and SNMP trap manager hosts can be configured per switch. An interconnect switch SNMP scripting utility and HP recommended example scripts are included.

This capability allows the interconnect switch to be monitored remotely from an SNMP/RMON network management station such as [Insight Manager 7](http://www.compaq.com/products/servers/management/cim7-description.html)³ or [HP OpenView](http://www.hp.com/products1/softwareproducts/software/openview/index.html).⁴ The interconnect switch may also be configured through the HP OpenView Network Node Manager. The unique interconnect switch integration with Insight Manager 7 provides the following additional features:

- Pre-registered (compiled) GbE Interconnect Switch MIBs.
- Automatic discovery and identification of the interconnect switches upon registration of MIBs.
- Automatic receiving of traps and events from the interconnect switches upon registration of MIBs.

³ Available at <http://www.compaq.com/products/servers/management/cim7-description.html>.

⁴ Available at <http://www.hp.com/products1/softwareproducts/software/openview/index.html>.

- Graphical representation of the p-Class server blades and interconnect switches.

Switch security

The GbE Interconnect Switch uses a layer 2 access control list or filtering database to segment the network, control communications between segments, and provide intrusion control. The switch allows manual entry of specific media access control (MAC) addresses to be filtered from the network. Filtering of both unicast and multicast traffic is possible. The maximum number of MAC addresses learned on a per port basis may further be restricted.


Several additional features are provided on the GbE Interconnect switch to allow the switch administrator to secure the management interfaces. These features include the ability to:

- Configure multiple password protected accounts with various levels of access.
- Specify the IP-based management stations that are allowed to access each switch.
- Set a user interface idle time-out period.
- Disable web-based and Telnet access.
- Configure port-based IEEE 802.1Q tagged VLANs for server grouping and data isolation.

Serviceability and diagnostics

The GbE Interconnect Switch has several serviceability features such as the front-mounted, hot-swappable interconnect switch blade, and the front panel RS-232 port, Ethernet ports, and a power re-set switch. The interconnect switch also has system status LEDs and a LED annunciation panel that displays the link speed and activity status for each port (Figure 7).

Figure 7. Front port LED panel for ProLiant BL p-Class GbE Interconnect Switch



*Uplink ports 21 and 22 on the p-Class C-GbE Interconnect Kit are 10/100/1000T. Uplink ports 21 and 22 on the p-Class F-GbE Interconnect Kit are 1000SX.

The optional HP Diagnostic Station (Figure 8) is available to configure and diagnose an interconnect switch or a server blade that has been removed from the server blade enclosure. The Diagnostic Station provides power to the interconnect switch and can be used to perform any of the following tasks:

- Build the initial interconnect switch configuration.
- Download a switch configuration file.
- Test network adapter and interconnect switch port activity.
- Diagnose interconnect switch status.
- Test an interconnect switch after installing an option or upgrade.

Figure 8. Diagnostic Station



The GbE Interconnect Switch provides many additional serviceability and diagnostic features including:

- Port mirroring with the ability to mirror desired type of frames (egress, ingress, or both).
- Power on self test (POST) at boot for hardware verification.
- Monitoring screens via the user interfaces for port utilization, data packets received/transmitted, error packets, packet size, trunk utilization, SNMP data, etc.
- Details of system information via the user interfaces such as port parameters and link status, switch asset information, configuration values, log entries, etc.
- The ability to "ping" to test the connectivity on the Ethernet network.
- Local system log (syslog) with ability to view and clear messages that may be saved (uploaded) as text file via TFTP.
- MAC addresses view, clear, and delete from the forwarding database for identifying problems with MAC address learning and packet forwarding.
- The ability to set the interconnect switch to a valid firmware image in case of firmware corruption.

For more detailed information on the administration capabilities of the ProLiant BL p-Class GbE Interconnect Switch, see the ProLiant BL p-Class GbE Interconnect Switch user guides.⁵

Virtual LAN

The GbE Interconnect Switch supports 256 port-based IEEE 802.1Q virtual LANs (VLANs) with GVRP dynamic VLAN registration⁶. Members of a VLAN may be untagged and tagged ports according to IEEE 802.3ac VLAN Ethernet frame extensions for 802.1Q tagging. Therefore, GbE Interconnect

⁵ Available at <http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html>.

⁶ GVRP stands for GARP VLAN Registration Protocol.

Switch VLANs may span other switches that support 802.1Q tagging located within the network infrastructure.

Spanning tree

The GbE Interconnect Switch meets the IEEE 802.1D spanning tree protocol (STP). Users can configure STP switch parameters, including priority and cost, on a per port basis. Each interconnect switch can automatically find the STP root bridge on the network. Otherwise, the interconnect switch will act as the root bridge for the STP domain.

For networks designed without loops or individual switch ports connected to server blades or other devices where a loop does not exist, the GbE Interconnect permits STP to be:

- Disabled on a per switch or per port basis.
- Placed in a bypass mode (Cisco PortFast equivalent) on a per port basis allowing a port to skip the standard STP modes and enable itself directly in the forwarding state.

EtherChannel compatible link aggregation

The GbE Interconnect Switch complies with IEEE 802.3ad static link aggregation (excluding LACP⁷) compatible with Cisco EtherChannel (Fast EtherChannel, Gigabit EtherChannel). Each GbE Interconnect Switch supports six multi-port trunks with up to eight ports per trunk.

Internet group management protocol

The GbE Interconnect Switch provides internet group management protocol (IGMP) snooping v1 and v2, configurable to a non-querier mode. The IGMP state may be enabled and disabled on a per VLAN basis as well as a configurable response report delay and query interval. The GbE Interconnect Switch allows a maximum of 191 concurrent multicast groups (127 dynamically learned by IGMP, 64 static multicast).

Data storm prevention

The GbE Interconnect Switch permits configurable thresholds (in packets per second) to prevent three types of packet storms: broadcast, multicast, and destination address (DA) unknown. If the threshold is exceeded, any additional packets received would be dropped.

Quality of service

Support for quality of service (QoS) IEEE 802.1p on the GbE Interconnect Switch allows switch administrators to set priority levels on each switch for forwarding packets. Each switch supports four classes of traffic (buffers or queues) for implementing priority based on the priority tag of the packet. Administrators can map up to eight priority levels to four classes. Traffic from a specific server port can be given priority over packets from other devices according to this range of priority levels. For example, with multiple packets in a buffer, the packet with the highest priority would be forwarded first, regardless of when it was received.

Enterprise-class performance

The GbE Interconnect Switch includes the following performance features:

- Nonblocking, full wire speed on all ports.
- 9.6 Gb/s external port bandwidth per server blade enclosure (full duplex).
- 7.2 million packets per second maximum external port frame forwarding throughput per server blade enclosure (64-byte packets).
- 8,191 MAC addresses per switch with automatic MAC address learning.

⁷ Link Aggregation Control Protocol.

- 32-MB Main, 6-MB flash, and 16-MB packet buffer memory per switch (packet buffer memory shared between ports).

ProLiant Network Adapter Teaming

ProLiant Network Adapter Teaming is typically used to provide fault resiliency for the NICs within the team. If one NIC should fail, the surviving NICs maintain network connectivity. However, network adapter teaming may also be used to maintain network services if an Ethernet switch fails. The architecture of the p-Class system allows for both network adapter and interconnect switch fault tolerance when using ProLiant Network Adapter Teaming.

The ProLiant p-Class servers include standard HP NC series network adapters that support all three types of ProLiant Network Adapter Teaming:

- Network Fault Tolerance (NFT)
- Transmit Load Balancing (TLB)
- Switch-assisted Load Balancing (SLB)

When teaming HP NC series network adapters with the interconnect switch, NFT and TLB teaming are the same as with any ProLiant system. However, for SLB teaming, the pre-defined switch-to-server architecture across the server blade enclosure's backplane must be considered. SLB requires that all NICs in the team must be connected to the same switch.

The GbE Interconnect Switch's redundant crosslinks add significant advantages when configuring teaming for switch fault tolerance. For details, see the [HP ProLiant network adapter teaming](#) white paper⁸ located on the ProLiant teaming home page.

Conclusion

The ProLiant BL p-Class system is a comprehensive solution for customers operating multi-tiered environments that demand space and power efficiency, on-demand computing, end-to-end system redundancy, rapid hardware and software deployment, and scalable management. The ProLiant p-Class GbE Interconnect Kits are part of the family of interconnect options available for the ProLiant BL p-Class system. Each GbE Interconnect Kit includes a pair of fully redundant, managed Ethernet switches that consolidate up to thirty-two network adapters (at a maximum of 10/100 Mb/s) to one to twelve external ports. These options are ideal for reducing Ethernet network cabling and the time required to deploy, manage, and service ProLiant BL p-Class systems.

⁸ Available at <http://ftp.compaq.com/pub/products/servers/networking/TeamingWP.pdf>.

Appendix: Summary of features of the ProLiant BL p-Class GbE Interconnect Kit

Switch performance

- Non-blocking full wire speed architecture
- Store and forward mode layer 2 switching standard
- Support for 3rd party external layer 7 content delivery switch / appliance
- Support for F5 BIG-IP Blade Controller load balancer and traffic management
- Frame forwarding throughput of 1.488 million packets per second (per Gigabit port)
- Frame forwarding throughput of 0.1488 million packets per second (per Fast Ethernet port)
- 7.2 million packets per second maximum external port (uplink) frame forwarding throughput per server blade enclosure (64-byte packets)
- 9.6 Gb/s External port (uplink) bandwidth per server blade enclosure (full duplex)
- 32 MB Main, 6 MB flash, and 16 MB packet buffer memory per switch (packet buffer memory shared between ports)
- Auto-negotiation and auto-sensing with full-duplex support and ability to manually force port speed and duplex mode
- Auto-MDI/MDIX on all ports enabled with auto-negotiation
- 8,191 MAC addresses per switch with automatic MAC address learning
- ARP for IP to MAC address resolution

Switch network features

- IEEE 802.3 10Base-T Ethernet, IEEE 802.3u 100Base-TX Ethernet, IEEE 802.3ab 1000Base-T Ethernet, and IEEE 802.3z 1000Base-SX
- IEEE 802.1D spanning tree protocol (mono-spanning tree)
- Spanning tree bypass fast forwarding mode on a per port basis (Cisco PortFast equivalent)
- Enable/disable and configure spanning tree port cost and priority on a per port basis
- IEEE 802.3ad link aggregation (excluding LACP) supporting up to 6 multilink trunk groups with 8 ports per group; compatible with Cisco EtherChannel trunking (Fast EtherChannel, Gigabit EtherChannel)
- 256 IEEE 802.1Q port based VLANs per switch
- IEEE 802.3ac VLAN Ethernet frame extensions for 802.1Q tagging on a per port basis
- Tagged and untagged ports as members of a VLAN
- GARP VLAN registration protocol (GVRP) providing 802.1Q compliant VLAN pruning and dynamic VLAN creation
- IEEE 802.1p QoS with 4 classes of service mapped to 8 priority levels
- IGMP snooping v1 and v2, configurable to a non-querier mode
- IGMP state enabling and disabling on a per VLAN basis
- IGMP response report delay and query interval configuring
- 191 Maximum concurrent multicast groups (127 dynamically learned by IGMP, 64 static multicast)
- Broadcast, multicast, and unknown packet storm control with a configurable threshold value
- IEEE 802.3x flow control with manual configuration capability
- Primary and secondary SNTP (simple network time protocol) server with time zone support and automatic daylight savings adjustment

Switch deployment and configuration

- Supports any combination of ProLiant BL p-Class server blades
- Default pre-configuration for immediate plug-in operation in the server blade enclosure
- Communicate to any and all server blade network adapters from any Ethernet external port
- Manage both switches, access all iLO ports, and execute PXE from any external Ethernet port
- Web-based interface accessible from any switch Ethernet port
- Menu driven console interface accessible from any switch port

- Command line interface (CLI) with scripting capability accessible from any switch port
- Telnet access to the CLI and menu-driven console interfaces accessible from any switch port
- SNMP-based scripting with best-case HP recommended example scripts
- Integrated switch scripting within Rapid Deployment Pack for Windows v1.40 or greater
- Altiris Switch Add-on utility in Rapid Deployment Pack for Windows v1.40 or greater
- One serial and two Ethernet ports conveniently located on the front of each switch for additional uplinks and local management, port mirroring, and other administration and diagnostic tasks
- Serial line IP (SLIP) access to command line and menu-driven console interfaces
- Configurable forwarding MAC address aging time settable to any value from 1 to 1 million seconds (default is 300 seconds)
- MAC address user management sorting on a per port and per VLAN basis
- Manual (static) entries in MAC address table
- Manual, or automatic IP settings via a DHCP or BOOTP server
- Ability to restore switch to factory default settings
- TFTP to upload and download (save, restore, and update) the switch configuration file
- TFTP to upload and download (save, restore, and update) the switch firmware
- ZModem to download the switch firmware (restore and update) via the serial interface
- Switch configuration retention after firmware upgrade
- HP Support Paq automated firmware upgrade process for Windows deployment stations
- Human read/write configuration file for viewing, printing, and editing
- Pre-configured customized port naming with respect to server blade NIC connectivity
- Per port bandwidth control of ingress and egress traffic
- Ability to name ports on a per port basis
- Full ability to enable and disable any port (both internal and external ports) on both switches

Switch diagnostics and monitoring

- One serial port and two Ethernet ports conveniently located on the front of each switch for additional uplinks and local management, port mirroring, and other administration and diagnostic tasks
- Front panel per port speed and per port link activity LED annunciation panel per switch
- System and management status LEDs
- Per port speed and link activity LEDs adjacent to all external Ethernet ports
- Active virtual graphic in the web-based switch interface
- Port mirroring with ability to mirror desired type of frames (egress, ingress, or both)
- Switch statistic monitoring including port utilization, data packets received/transmitted, port error packets, packet size, trunk utilization, SNMP data, etc
- System reporting such as port parameters and link status, switch asset information, configuration values, log entries, etc.
- Ping capability to test the connectivity on the Ethernet network
- SNMP v1 with four configurable community strings and SNMP trap manager hosts
- MIB-II, Bridge MIB, Interface MIB, Extended Bridge MIB, Ethernet-like MIB, Entity MIB, and HP enterprise switch MIBs
- Bridge, remote monitoring, and switch environmental traps
- Pre-registered (compiled) switch MIBs with Insight Manager 7
- Insight Manager 7 automatic discovery, identification, and receiving of traps and events
- Insight Manager 7 graphical representations of switches
- Insight Manager 7 standard database activities (queries, tasks, reporting, etc)
- Power on self test (POST) at boot for hardware verification
- Portable Diagnostic Station to configure, upgrade, and diagnose an interconnect switch and server blade removed from the rack environment
- Ability to return switch to a valid firmware image in case of firmware corruption
- Local system log (syslog) with ability to view and clear messages, and save (upload) as text file via TFTP

Switch security

- Password protected multi-level user accounts supported on all management interfaces
- Configurable user interface idle time-out period
- Ability to disable web-based and Telnet access to the switch user interfaces
- 256 Port-based IEEE 802.1Q tagged VLANs per switch (512 per server blade enclosure)
- Ability to specify the IP-based management stations that is allowed to access the switch
- Unicast and multicast static MAC address packet filtering table (layer 2 access control list)
- Restriction on the maximum number of MAC addresses learned on a per port basis

Switch availability

- Front-mounted, hot-swappable switch with configuration retention and no need to cable/recable when removing
- Redundant switches per server blade enclosure
- Redundant uplink ports per switch
- Redundant front panel management/diagnostic ports per switch
- Redundant crosslinks for switch to switch communication and failover scenarios within the server blade enclosure
- Redundant pairs of network adapters per server routed to different switches
- Redundant N+N hot-plug redundant power to each switch
- Redundant N+1 on-board cooling per switch
- IEEE 802.3ad automatic multi-link load balancing and link failover (excluding LACP)
- Load balancing of unicast traffic
- ProLiant network adapter teaming
- Redundant configurable community strings and SNMP trap manager hosts
- Redundant configurable SNMP servers

Switch ports per server blade enclosure

- Four external 10/100T Fast Ethernet ports and four external 10/100/1000T (C-GbE) or 1000SX (F-GbE) Gigabit Ethernet ports, all on a pair of hot-pluggable rear-mounted interconnect modules
- Four external 10/100T Fast Ethernet ports on switch front panels
- Two external DB-9 serial ports on switch front panels (one per switch)
- 32 internal 10/100 Fast Ethernet ports to server blade network adapters (4 per server bay)
- Two internal 10/100 Gigabit Ethernet inter-switch crosslinks for switch to switch communication and failover scenarios
- I²C Switch to management module communications
- All external Ethernet ports may be used for data, switch and iLO management, and/or PXE remote configuration.
- All internal Ethernet signals routed as Ethernet across individual CAT5e signal traces
- 12 RJ-45 (C-GbE) or 4 LC fiber + 8 RJ-45 (F-GbE) external Ethernet port connectors

Switch physical and environmental

- AC power input (at server blade system level)
- Direct facility -48 DC power input
- 40 Typical and 50 maximum power consumption per switch (watts)
- Local front power reset/power cycle button on each switch
- Remote reset/power cycle via user interface
- 10 to 35 operating and -30 to 60 storage temperature (Celsius)
- 20% to 80% operating and 5% to 95% storage relative humidity
- FCC Class A, ICES-003 Class A, AS/NZS 3548 Class A, and VCCI Class A electromagnetic interference (EMI) certifications
- UL/CUL and CE safety certifications
- 11.5 pounds per interconnect switch and interconnect module

For more information

For additional information, refer to the resources detailed below.

Resource description	Web address
<i>ProLiant BL p-Class GbE Interconnect switch home page</i>	http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html
<i>ProLiant BL p-Class Networking Overview white paper</i>	http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html
<i>Using ProLiant Essentials Rapid Deployment Pack for scripted blade-based switch configuration white paper</i>	http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html
ProLiant BL p-Class GbE Interconnect Switch user guides	http://h18004.www1.hp.com/products/servers/proliant-bl/p-class/bl-p-interconnect-switch.html
ProLiant network adapter teaming	http://www.compaq.com/products/servers/networking/teaming.html
F5 BIG-IP Blade Controller for ProLiant BL systems	http://h71028.www7.hp.com/enterprise/html/4557-0-0-121.html

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

BIG-IP is a registered trademark of F5 Networks, Inc. All other trademarks or registered trademarks mentioned in this document are the property of their respective owners.

5982-2729EN, 05/2004

